

T3 Network Nears Full Production

A year of dedication and hard work by the NSFNET partnership has begun to pay dividends for network users as the ANS T3 technology becomes a production tool for NSFNET sites.

As of March 1, a number of midlevel networks were moving their traffic across the T3 network and plans were in place to move the remaining networks over the next few months.

Merit obtains NSFNET backbone services from ANS which provides a major national network that operates at T3 speeds using circuits provided by MCI and central networking technology based on the IBM RS/6000™.

Fine-tuning over the past year

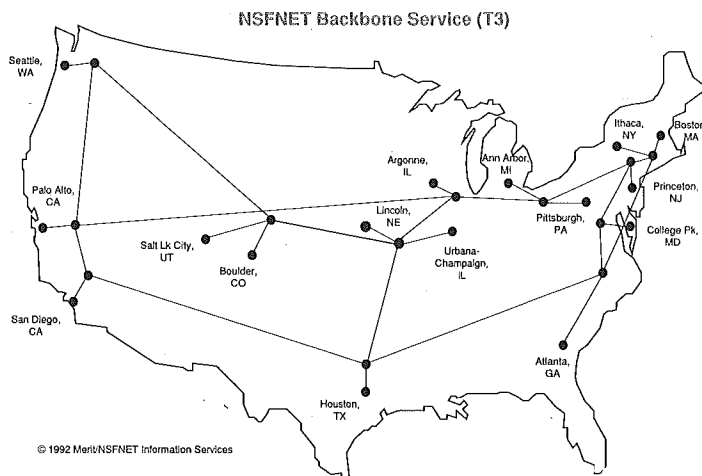
Activities over the past twelve months have focused on installing T3 technology at all NSFNET sites, re-engineering the early T3 backbone to expand total bandwidth and provide greater redundancy, and refining the initial developmental hardware and software to bring the level of reliability up to production standards. In addition, some organizational changes have been made to ensure a smooth problem-solving path.

The first stage is nearly complete and a plan is in place to deploy higher speed T3 interface cards which will increase the throughput on all nodes during second quarter 1992.

Strengthening the architecture

When T3 connections were originally implemented at the beginning of 1991, some problems were experienced with the new ar-

chitecture. This led to a major effort by Merit, ANS, IBM, and MCI to improve the technology before T3 traffic was increased. Debugging activities continued at a fast pace



through October 1991, until a reasonable degree of reliability was achieved.

Peeling the onion

The problem analysis and resolution process has been compared to "peeling the skin of an onion", as multiple glitches that at first appeared to have the same symptoms were

See T3 production, page 8

INSIDE:

2—Merit seminar

3—WAIS

5—Inside the T3 network

10—NSF solicitation

11—NSF acceptable use policy

12—Tooling up with Rover

March/April 1992 — Vol. 5 No. 1



Odds Favor the User at Upcoming Merit Seminar!

Do you want to learn more about the Internet and the many resources available through NSFNET backbone services? Your opportunity is at hand, as Merit/NSFNET provides another high-quality, information-packed day-and-a-half seminar in Las Vegas, Nevada, on June 1 and 2, 1992.

The revolution in national computer networking continues to affect the way people collaborate in education and research. Hundreds of thousands of educators and researchers are accessing library resources, supercomputers, video applications and the expertise of colleagues around the world through their connections to NSFNET and the Internet.

Who Should Attend:

This informative seminar is intended for providers of campus information and administrators as well as beginning and intermediate network users.

Facilities

The seminar will be held on the campus of the University of Nevada, Las Vegas, in a state-of-the-art facility. Attendees will be able to access remote facilities through NevadaNet and NSFNET.

Lodging

Lodging will be available at the St. Tropez, a non-gaming hotel which is reserving a block of rooms for seminar attendees until May 1, 1992. Make your own room reservations and payment arrangements directly by calling 800-666-5400 or 702-369-1150.

Travel discounts

Special discounted airfares with Northwest Airlines are available for this seminar.

Save with early registration

The registration fee is \$295 until May 15, 1992. After May 15 the cost is \$345. Limited scholarships are available; contact us today.

The fee includes the one-and-a-half-day seminar, receptions on Sunday and Monday evenings, lunches on Monday and Tuesday, refreshments, the demonstration room, and all seminar materials. For more information on any or all of the above, you may reach us electronically at: seminar@merit.edu or by phone at: **1-800-66-MERIT** or **313-936-3000**.

— Merit/NSFNET Information Services

This informative seminar is intended for providers of campus information and administrators, as well as beginning and intermediate network users.

The *Link Letter* is published by:
Merit/NSFNET Information Services
2901 Hubbard, Pod G
Ann Arbor, Michigan 48109-2016

To subscribe to the *Link Letter* send an electronic message to:
NSFNET-Linkletter-Request@merit.edu

For further information, contact
Patricia Smith, Editor, at 1-800-66-MERIT or send e-mail to:
NSFNET-info@merit.edu

Suggestions and contributions welcome.
Permission to reprint all or part of the *Link Letter* for non-profit purposes is granted, provided full acknowledgement of the source, including authors' names, is given.



This material is based upon work supported by the National Science Foundation under Grant No. NCR 8720904. The Government has certain rights in this material. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

WAIS—Making it Easier to Access Internet Resources

(Reprinted from *CERFnet News*, Volume 3 Number 6. Used with permission.)

The Wide Area Information Servers (WAIS) system is a set of products supplied by different vendors to help end-users find and retrieve information over networks. Using WAIS (pronounced ways) the user isn't required to become familiar with several different systems. Also, since the interface consolidates data from many sources, the data can be manipulated effortlessly, virtually without regard to their origins.

Thinking Machines, Apple Computer, and Dow Jones initially implemented such a system for use by business executives. These products are becoming more widely available from various companies.

What does WAIS do?

Users on different platforms can access personal, company, and published information from one interface. The information can be anything: text, pictures, voice, or formatted documents. Since a single computer-to-computer protocol is used, information can be stored anywhere on dif-

ferent types of machines. Anyone can use this system since it uses natural language questions to find relevant documents. Relevant documents can be fed back to a server to refine the search. This avoids complicated query languages and vendor specific systems. Successful searches can be automatically run to alert the user when new information becomes available.

How does WAIS work?

The servers take a user's question and do their best to find relevant documents. The servers, at this point, do not understand the user's English language question, rather they try to find documents that contain those words and phrases and ranks them based on heuristics. The user interfaces (clients) talk to the servers using an extension to a standard protocol Z39.50. Using a public standard allows vendors to compete with each other, while bypassing the usual proprietary protocol period that slows development. Thinking Machines is giving away an implementation of this standard to help vendors develop clients and servers.

See WAIS, following page

Users on different platforms can access personal, company, and published information from one interface.

WAIS—Mac Interface Proposed

The latest and greatest from the world of WAIS was demonstrated for attendees at a WAIS seminar held at the MCNC Center for Communications in Research Triangle Park, NC on February 3 and 4.

Participants viewed a number of demonstrations of WAIS software on several platforms and were treated to a sneak preview of an Apple development project to create a new Macintosh-WAIS interface. The prototype of this interface combines

graphics and text to present information in a newspaper-like "front page" format.

Speakers at the seminar included meeting organizer George Brett of MCNC, Cliff Lynch of the University of California, Brewster Kahle of Thinking Machines, and Jim Fullton of University of North Carolina—Chapel Hill.

Further information about WAIS may be obtained by contacting George Brett, ghb@concert.net.

—Mark Davis-Craig, Merit/MichNet

"These connections allow scientists to access the very large space and earth science data bases at the Goddard and Ames Research Centers."

Two NASA Sites Connected

Merit has received additional funding under the National Science Foundation cooperative agreement to provide T3 connectivity to two NASA sites, the NASA-Ames Research Center and NASA-Goddard Space Flight Center. These connections, one directly to Ames and a second for Goddard at College Park, were expected to be in place by early March 1992.

"These connections allow scientists to access the very large space and earth science

data bases and computational facilities at the Goddard and Ames Research Centers to achieve major breakthroughs in large-scale modelling, simulations, and space mission planning," stated Tony Villasenor, Program Manager for the NASA Science Network. "This is another example of the fine collaboration continuing between the NASA and NSF communities."

—Merit/NSFNET Information Services

WAIS, cont. from page 3

Where do WAIS servers exist?

Even though the system is very new, there are already several servers:

- Dow Jones is putting a server on their own DowVision network. This server contains the Wall Street Journal, Barons, and 450 magazines. This is a for-pay server.

- Thinking Machines operates a Connection Machine on the Internet for free use. The databases it supports are some patents, a collection of molecular biology abstracts, a cookbook, and the CIA World Factbook.

- MIT supports a poetry server with a great deal of classical and modern poetry.

- Cosmic is serving descriptions of government software packages.

- The Library of Congress has plans to make their catalog available on the protocol.

- Weather maps and forecasts are made available by Thinking Machines as a repackaging of existing information.

- The "directory of servers" facility is operated by Thinking Machines so that new servers can be easily registered as either for-

pay or for-free servers and users can find out about these services.

For more information

Contact Brewster Kahle for more information on the WAIS project, the Connection Machine WAIS system, or the free Mac, Unix Server, and X Window System interfaces. There is a mailing list that has weekly postings on progress and new releases; to subscribe send e-mail to: wais-discussion-request@think.com.

You can also retrieve a copy of the WAIS Bibliography by Barbara Lincoln, Thinking Machines, October, 1991. It is available via anonymous FTP from [quake.think.com](ftp://quake.think.com). The directory/file is: [/pub/wais/wais-discussion/bibliography.txt](ftp://pub/wais/wais-discussion/bibliography.txt). It is also available via the WAIS server [wais-discussion-archive.src](ftp://wais-discussion-archive.src).

Brewster Kahle is the Project Leader for Wide Area Information Servers.

Kahle's Internet-mail address is brewster@think.com.

— Brewster Kahle

Thinking Machines operates a Connection Machine on the Internet for free use.

The T3 Network—What Makes It Tick?

As the migration of all production traffic to the ANS T3 backbone nears completion, performance improvement is the feature that will be most apparent to a majority of users of NSFNET backbone services. However, underlying the visible product is a complex foundation of innovations and advances which reside at the forefront of network technology.

Architecture

The architecture for the T3 network evolved from many observations made in the course of managing the T1 network and is shown in the figure at right.

Packet switches in the T3 infrastructure are designated as either Exterior Nodal Switching Subsystems (ENSSs) or Core Nodal Switching Subsystems (CNSSs). ENSSs are located on the regional network premises and CNSSs are co-located at MCI carrier switching centers which are also known as "points-of-presence" (POPs) or "junction points." The decision to co-locate packet switches within POPs was key for several reasons. Since these locations are major carrier circuit switching centers, they have the full backup power essential to the stability of the network and many are staffed 24 hours/day, 7 days/week.

Additionally, co-locating the CNSSs provides the opportunity to more closely match the carrier-provided circuit switched network topology with packet switched network topology. As a result, path redundancy becomes much easier to engineer.

With this architecture, the need for transit traffic to traverse the end nodes is eliminated. Only traffic specifically sourced from or destined to a given ENSS traverses the access link. Consequently, the total bandwidth needs of the CNSS cloud are

based on aggregate requirements because of the shared portion of the network while access link needs are based on individual requirements.

Regional network connections

Currently the T3 infrastructure supports both Ethernet and FDDI interconnection. The access into the CNSS cloud is via DS3 (45 Mbs), DS1 (1.544 Mbs) or DS0 (56Kbs) circuits. The network uses cisco™ packet switches for 56 Kbs access into the backbone.

DS3 circuits

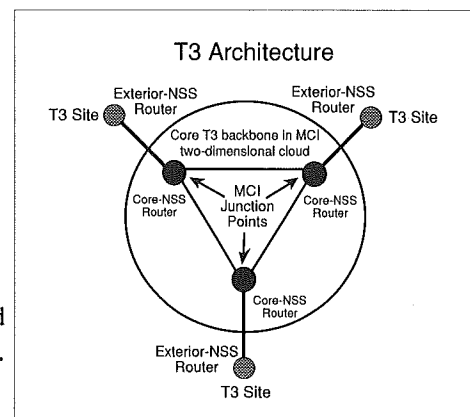
The DS3 circuits operate in non-subrated mode, which means that the full 45 Mbs, less framing and carrier management overhead, are available for user traffic. The physical and electrical interfacing to these lines is handled by a Data Service Unit (DSU). Nearly all of the DS3 circuits are carried on terrestrial fiber-optic lines.

Packet switches

The T3 network packet switches are based on the IBM RISC/6000™ machines running code based on the AIX operating system, with appropriate network adapter hardware, and packet switching, routing and network management software. In order to drive the DS3 lines at 45 Mbs, high performance adapter cards and software were required.

To accomplish this, "smart" adapter cards with the Intel 960CA 32-bit microprocessor on board were developed. The cards contain all information required to switch a packet, including routing tables and stored

See T3 Tech, following page



T3 Tech, cont. from page 5

In fact, "smart cards" will switch packets between themselves without any main system intervention.

instruction code for packet switching. In fact, "smart cards" will switch packets between themselves without any main system intervention, thereby greatly increasing packet switching efficiency. CPU intervention is necessary only when routing information is updated. In addition, the presence of a high speed system bus (IBM's Microchannel™) ensures that bus bandwidth will not restrict packet flow. Deployment of these cards at the POP-based CNSS and ENSS sites is being finalized. Completion of this task is expected to take several months.

Routing

The routing technology and architecture used on the T3 network is essentially the same as the T1 network. Routing is fully distributed, with each NSS exchanging state information with other nodes and then computing the optimal routes. The Intermediate System-to-Intermediate System (IS-IS) link state routing protocol is used as the intra-domain routing protocol.

In order to take advantage of the new architecture and to minimize routing overhead, exterior information is transmitted internally via the Internal Border Gateway Protocol (IBGP). With over 4,500 network numbers in the routing tables maximizing the efficiency and controlling the dynamics of this information transmission assumes paramount importance.

Border Gateway Protocol (BGP) or Exterior Gateway Protocol (EGP) are used as the inter-domain routing protocol, to share reachability and routing information with the attached regional networks. Policy-based routing with tightly controlled routing information flow between attached networks ensures that the T3 infrastructure will continue to serve as the primary transit WAN in the global Internet system.

In order to take advantage of the new architecture exterior information is transmitted internally via the Internal Border Gateway Protocol (IBGP).

T1 and T3 co-existence

Since the T3 network was designed as an overlay to complement the existing T1 network and eventually replace it, controlled traffic migration from the T1 to the T3 network posed some challenges.

In the migration phase, which is currently in progress, some sites are connected to both networks, and some to only one. It was essential that a stable and redundant cross-connect system be designed to allow user traffic to reach any site, regardless of how the destination site was connected or where the packet entered the network. Regional networks that were multihomed to both T1 and T3 will maintain peering with both backbones. This will allow the T1 network to be used to connect to regionals connected directly to the T1, and the T3 network to be used to connect to regionals directly connected to the T3. At press time traffic migration was underway and a load sharing configuration was running on the production network.

T1/T3 Interconnect Gateways

The routing design for load sharing involves T1/T3 interconnect gateways at Ann Arbor, Houston, San Diego, and Princeton with 5th and 6th gateways planned at Denver and Seattle. This scheme is integral to maintaining traffic balance and minimizing congestion at the interconnect points.

As 1992 progresses, developers and engineers within the partnership of ANS, Merit, IBM, and MCI will continue to provide a leading-edge environment for the growing number of network users.

—Bilal Chinoy, San Diego Supercomputer Center, and Pat Smith, Merit

Merit Proposals Accepted by CNI

Merit has received word that two proposals, one for an Internet directory project called TopNode, and the other for an Internet Information Packet project, have been accepted by the Coalition for Networked Information (CNI). CNI is a consortium formed by American Research Libraries, CAUSE, and EDUCOM to promote the creation of and access to information resources in networked environments in order to enrich scholarship and to enhance intellectual productivity.

Merit's responsibilities in the TopNode project include sharing knowledge about network resources, assisting in the gathering of such information, and creating an X.500 version of the TopNode directory.

In the other endeavor, Merit will work with CNI to create information packets about networking and the Internet. Participation in this project includes serving as a liaison among the various national networking organizations, the regional networks, and other organizations creating information packets about the Internet. Merit's focus will be on coordination and cooperation rather than re-writing the necessary materials.

If you have any questions about these projects, please contact Laura Kelleher, CNI Project Coordinator, Merit Network at 313-936-3000 or via email at lak@merit.edu

— Laura Kelleher, Merit/NSFNET

Merit's responsibilities in the TopNode project include sharing knowledge about network resources and creating an X.500 version of the TopNode directory.

High-Performance Computing Bill is Law

On December 9, 1991, the High Performance Computing Bill of 1991, known in its various iterations as "The Gore Bill" or "The NREN Bill," was signed into law by President Bush.

"Congressional action on the High Performance Computing and Communications Bill provides strong national recognition of the importance of the national networking effort in which Merit, Inc., has played a pioneering role," commented Dr. Douglas Van Houweling, Merit Board of Directors member and Vice-Provost for Information Technology at the University of Michigan. "The hard work is just now beginning, however, for the ex-

pectations of our colleagues and the competition for access are continuing to rise rapidly. Merit will continue to work with the networking community in response to these demands."

Obtaining an electronic copy

Readers may retrieve an electronic copy of the document from the Merit/NSFNET Information Services server, nis.nsf.net, via file transfer protocol (FTP). The file to retrieve is [/internet/legislative.actions/nrenbill.txt](ftp://internet/legislative.actions/nrenbill.txt)

— Merit/NSFNET Information Services

"The hard work is just now beginning, however, for the expectations of our colleagues and the competition for access are continuing to rise rapidly."

New organizational structures have been developed to enhance network operation and improve the interaction with regionals on network problems.

T3 production, cont. from page 1

discovered and fixed. Problems were investigated starting at the circuit level, moving to the DSU or interface hardware, and finally to the router software and hardware.

In recent months, T3 circuit and router deployment has continued at the sixteen NSFNET regional sites with the last of these completed in November. Nevertheless, the phase-in of additional production traffic to the T3 network was delayed because of continuing concerns over reliability and performance of the new technology. These issues have been addressed through a series of actions including:

- 1) improve the network software and hardware,
- 2) deploy new monitoring tools to better track problems,
- 3) provide a fallback system in the event of core node isolation, and
- 4) develop new test strategies to ensure that the final technology implemented would meet the standards of reliability to which NSFNET users had become accustomed.

In addition, new organizational structures have been developed to enhance network operation and improve the interaction with regionals on network problems.

Visible improvements

At the top of the stack of visible improvements are performance capabilities which now exceed that of most local networks to which users are attached. The national infrastructure comprises local/campus area networks (local area networks or LANs) connected to the regional and national networks (wide area networks or WANs).

In the past, constraints in performance were generally related to the WANs having slower speeds than the LANs. Many LANs which interconnect to the regional and

national networks are Ethernets capable of 10 Mb/sec, much higher than the maximum 1.544 Mb/sec of the T1 networks to which the LANs are attached. In some cases, attachments to campus and other organization LANs have been at speeds of 56 Kbps or less.

Bandwidth limitation now at the local level

Implementation of the T3 infrastructure has produced a setting in which backbone capacity is higher than much of the interconnecting network bandwidth beneath it. As a result, bandwidth limitation now occurs for most sites at the midlevel/regional level rather than at the national network level. While 100 Mb/sec Fiber Distributed Data Interface (FDDI) is supported at some local and regional networks, most have yet to deploy this technology.

Stability period used for tests

A two-week hiatus in December was declared a period of "no change" on the T3 network. The time was used to judge the stability and reliability of the network in a steady state condition, to prepare for the cutover of T1 traffic, to test the T1/T3 interconnect gateway backup, and to perform tests which were scheduled and coordinated with the regional networks. The testing period proved successful and enhanced confidence in the network.

Safety Net

Safety Net has been put in place as a fallback in case all T3 paths to a core node become unusable. It represents the addition of 12 T1 links to interconnect with the T3 backbone Core Nodal Switching Subsystem (CNSS) nodes. These safety net links are in-

See T3 production, following page

T3 production, cont. from page 8

stalled between the MCI switching centers and do not connect to the Exterior Nodal Switching Subsystem (ENSS) nodes. The T1 link metrics are designed so that a T1 path is used only if all other T3 paths to adjacent CNSS nodes become unreachable.

Restructuring of Merit/ANS NOC and Internet Engineering Group

The NSFNET is a vital communications link which demands a trouble-shooting structure that is as efficient and reliable as possible. In an effort to produce such a framework the Network Operations Center and Internet Engineering groups have been modified to provide a three-tier problem-resolution setting.

1st Level: Network Operations Technicians. NOC operators will continue to provide first level technical support for network problems. The NOC function will become more formalized with emphasis on reporting, escalation, tracking, and procedure execution.

2nd level: National Network Attack Force. As part of the ongoing efforts to ensure complete trouble-shooting coverage of the national network, the National Network Attack Force (NNAF) has been formed. This group of six individuals provides 24-hours per day/seven-days per week implementation of the T3 network problem resolution process. The team develops NOC diagnostic tools, trains NOC operators, and performs network-related tests. The group reports to the IE manager.

3rd Level: Internet Engineering Group. A number of the tasks indicated above were previously handled by the Internet Engineering staff. As the NNAF members assume those duties, the Internet Engineers will have time to pursue longer-term engineering

activities and for attending to the highest priority problems which reach the third level of escalation.

Future T3

The mix of a solid foundation plus future leading edge technology such as the imminent deployment of the new IBM-developed, high-performance interfaces for the RS/6000-based routers, helps ensure that notable network performance will be realized.

As noted on page 10 of this issue, the National Science Board recently approved a request by the National Science Foundation to extend the current Cooperative Agreement with Merit for a period of up to 18 months beyond the November 1992 expiration date of the current contract.

Developers and engineers within the partnership continue to work toward a greater understanding of the interactions of application and network engineering, and to tune the network to maximize efficiency as the user base grows.

— Ellen Hoffman, Mark Knopper, and Pat Smith Merit/NSFNET

The mix of a solid foundation plus future leading edge technology helps ensure that notable network performance will be realized.

NSF Plans for Future Backbone Services

At the end of November 1992, the existing five-year cooperative agreement which the National Science Foundation has with Merit for NSFNET backbone services comes to an end. For those who have come to rely on the NSFNET for the work they do in research and education, the question of "what happens next?" is critical to assuring stability and planning for the future.

Some answers are at hand since the National Science Board met at the end of November and approved a plan put forth by NSF's Division of Networking and Communications Research and Infrastructure.

"NSF plans to make a draft solicitation for the Backbone follow-on available sometime in March for a two to three month period of public comment," said Steve Wolff. He added that NSF hopes to issue the solicitation in May or June of this year.

No service interruption for users

The plan has two aspects: the first is a provision which will assure that users will see no interruption in service when the current agreement expires. The second is a set of solicitations for longer-term provision of services.

"As an association of network and information service providers, FARNET [Federation of American Research Networks] has endorsed the principles of competition and choice, especially where the market for these services is maturing," said Laura Breeden, Executive Director of FARNET. "We are looking forward to working with NSF on the implementation of its plans for the INREN [Interim National Research and Education Network]."

Plan reflects FARNET recommendations

This past November FARNET submitted a number of recommendations to NSF regarding interregional connectivity after November 1992. The plan reflects some aspects of the FARNET input, as well as the input of many other individuals, federal agencies, and organizations with whom NSF consulted.

A complete copy of the document may be obtained via anonymous FTP to [nis.nsf.net](ftp://nis.nsf.net). The directory is [/nsfnet/news.releases/](ftp://nis.nsf.net/nsfnet/news.releases/) and the document to retrieve is:

[nsfnet.project.development.plan](ftp://nis.nsf.net/nsfnet/project.development.plan)

—*Merit/NSFNET Information Services*

New Acceptable Use Policy from NSFNET

The following is a new acceptable use policy for the NSFNET backbone service. Copies are available via anonymous FTP from **nis.nsf.net**. The file to retrieve is: **acceptable.use.policies/nsfnet.txt** or **nsfnet.ps** for the PostScript version.

The NSFNET Backbone Service Acceptable Use Policy

GENERAL PRINCIPLE:

(1) NSFNET Backbone services are provided to support open research and education in and among US research and instructional institutions, plus research arms of for-profit firms when engaged in open scholarly communication and research. Use for other purposes is not acceptable.

SPECIFICALLY ACCEPTABLE USES:

(2) Communication with foreign researchers and educators in connection with research or instruction, as long as any network that the foreign user employs for such communication provides reciprocal access to US researchers and educators.

(3) Communication and exchange for professional development, to maintain currency, or to debate issues in a field or subfield of knowledge.

(4) Use for disciplinary-society, university-association, government-advisory, or standards activities related to the user's research and instructional activities.

(5) Use in applying for or administering grants or contracts for research or instruction, but not for other fundraising or public relations activities.

(6) Any other administrative communications or activities in direct support of research and instruction.

(7) Announcements of new products or services for use in research or instruction, but not advertising of any kind.

(8) Any traffic originating from a network of another member agency of the Federal Networking Council if the traffic meets the acceptable use policy of that agency.

(9) Communication incidental to otherwise acceptable use, except for illegal or specifically unacceptable use.

UNACCEPTABLE USES

(10) Use for for-profit activities (consulting for pay, sales or administration of campus stores, sale of tickets to sports events, and so on) or use by for-profit institutions unless covered by the General Principle or as a specifically acceptable use.

(11) Extensive use for private or personal business. This statement applies to use of the the NSFNET Backbone only. NSF expects that connecting networks will formulate their own use policies. The NSF Division of Networking and Communications Research and Infrastructure will resolve any questions about this Policy or its interpretation.

Rover—Tooling Up for the Expanding Internet

Owing to the leading-edge nature of the NSFNET technology, appropriate monitoring and diagnostic tools are generally not commercially available.

Getting the “right tool for the job” no longer applies only to obtaining the latest battery-powered gizmo from the local “Hammers R Us” hardware store. The value of owning the proper tool holds true for the “net ops” who manage today’s rapidly growing computer networks as much as for carpenters, electricians or auto mechanics.

Since 1988, Merit has operated a Network Operations Center (NOC) in Ann Arbor, MI, for managing the NSFNET. Network capacity has expanded more than 700 times and network connections have grown from 170 in July 1988 to more than 4,500 in January 1992. New nodes and links are now being added almost weekly and data traffic has increased from 195 million to over 13 billion packets per month.

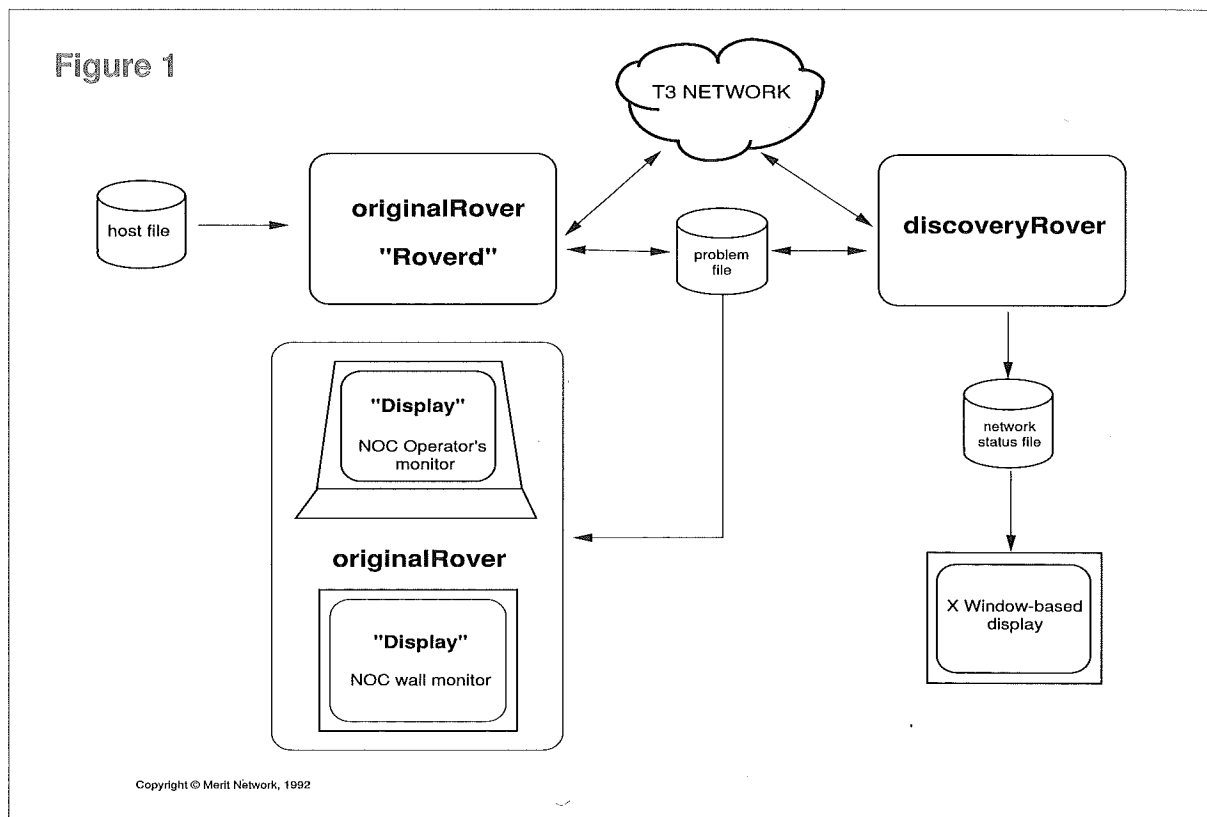
It is the job of the technicians in the Merit NOC to monitor this vast network, open trouble tickets for problems that occur, and track them through to closure.

Owing to the leading-edge nature of the NSFNET technology, appropriate monitoring and diagnostic tools are generally not commercially available. Consequently, Merit Internet engineers often write software to fill the need for timely problem detection and resolution on the backbone.

Rover is one of the primary tools used by the Merit NOC to monitor the NSFNET. The Rover package was written by Bill Norton of Merit’s Network Management Systems group.

“Bill has had almost daily contact with NOC personnel and has filled in as an

See Rover, following page



Rover, cont. from page 12

operator on occasion. The unique insight into the needs of the NOC gained from this contact is apparent in the design of the tool," said John Labbe, NOC assistant manager.

Rover is an extensible tool that currently incorporates two modules: originalRover, including Roverd and Display, and discoveryRover, which provides an in-depth, real-time picture of the Advanced Network & Services T3 network. Figure 1, at left, illustrates the interactions of the various pieces of Rover.

originalRover

originalRover was developed in 1988. It consists of two parts, Roverd and Display. Roverd reads information from a configuration file called **hostfile** that contains a list of network nodes, IP addresses and associated tests to be performed on each of them. Example lines from hostfile are shown in Figure 2, below.

If the test fails, a problem entry is added to the problem file. Subsequently, if a test succeeds the problem entry is removed. All problem additions and deletions are logged. The contents of the problem file are viewable on monitors through the Display module of originalRover.

Display

The Display module of originalRover allows the contents of the problem file to be viewed on four-by-five-foot monitors in the Merit NOC as well as on the NOC Ops' monitors. The problem file is periodically polled for modifications and the dis-

plays are updated accordingly. A representation of the text display is shown in Figure 3 on the following page.

Latest Developments—discoveryRover

During the past months, the T3 ANSnet (Advanced Network & Services network) has come into its own as noted in the article on page 1 of this issue. Rover has been made more robust to deal with the new technology and the increasing complexity of the networks.

Collection mechanism

The initial version of originalRover would cycle through the host file sequentially and the time it took to query all the nodes was somewhat dependent on the state of the network. It would not put a report of "node down" in the problem file until the reachability test—five successive "pings" (five retries, one second apart) had failed. This, plus the sequential pattern, allowed for a potentially high cycle time which was unacceptable.

With discoveryRover it is possible to obtain a more in-depth picture of the network by supplementing the original ping/reachability tests with Simple Network Management Protocol (SNMP) queries.

See Rover, following page

Figure 2

#	Nodename	IP Address	Info	NodeUp?	NetMgmtProto
	NSF_10_BGP-601-1_1	192.68.53.2	Help/CAEGP.HLP	PING()	
	NSF_5->NSF_10	129.140.5.11	Help/LLRT.HLP	PING()	LL(LL/NSS5-NSS10)
	NSF_5->NSF_8	129.140.5.12	Help/LLRT.HLP	PING()	LL(LL/NSS5-NSS8)
	NSF_5->NSF_12	129.140.5.13	Help/LLRT.HLP	PING()	LL(LL/NSS5-NSS12)
	NSF_5->NSF_18	129.140.5.14	Help/LLRT.HLP	PING()	LL(LL/NSS5-NSS18)
	NSF_5_spare	129.140.5.15	Help/SPR.HLP	PING()	
	NSF_5_NEARnet	129.140.5.16	Help/GW.HLP	PING()	
	NSF_5_NNStat	129.140.5.17	Help/NNSTAT.HLP	PING()	

Rover, cont. from page 13

Figure 3

```
originalRover Display v3.32      NOC Internet Rover      Wed Feb 12 17:23 EST
Cycle Time: 01:58 00:23          Wed Feb 12 22:23 GMT
                                   SSN[ 8]

# Reported      Problem      StatusLine

01 Wed_17:08 NOC_merit.edu 35.1.1.42 TROUBLE
02 Wed_06:14 NSF_10_AS-601 AS_601_Address ASDOWN tt 18512
03 Tue_17:18 ISIS cnss40-cnss32 DOWN 18504
04 Tue_17:18 ISIS cnss40-cnss41 DOWN 18504
05 Sat_15:59 NSF_4_Test 129.140.4.1 PING called, no answer
06 Sat_15:59 NSF_2_Test 129.140.2.1 PING called, no answer
07 Fri_13:03 AS_enss136-AS701(Alternet_Network_) DOWtt 18418
08 Fri_16:21 T3_E150_BGP-81-3/3 192.101.21.254 PING disabled by site
```

30-second discovery of the current topology

Only 30 seconds are needed for discoveryRover to determine the current topology of the network and store the state of the nodes and links in a network status file. This is accomplished by querying the “known” nodes for information on the nodes they know about, i.e. “what links do you have, who is attached to them, and what is the state of the link?” More specifically, SNMP queries return information on node reachability, link state information and information on Autonomous System (AS) reachability.

Nodes that respond to the query are marked up. Link state is determined by querying the IS-IS tables maintained at each of the nodes.

The IS-IS protocol operating over the links and the routers will mark the links **down** if they have not heard a hello packet recently. Conversely, if the router says the link is up, that means it has heard a hello packet and the remote machine is probably up which results in the node being marked busy—most

likely up but has not answered the SNMP query yet. Therefore, if SNMP link queries obtain a response then the node at the end of the link is assumed to be functioning and its status is marked **up**.

This discovery process is the mechanism that obtains the information stored in the Network Status File which contains a list of objects named nodes and links. Query responses may be displayed as ordinary text or through an X Window-based environment.

X Window-based display

An X Window-based graphical interface has been developed to aid in using Rover. This “front end” to the Rover data collector displays nodes and links in a color-key fashion which reflects their state as seen in the Network Status File. As new node and link objects are discovered in the network, the map is updated to show the new objects. Placement of nodes on the display is accomplished using the mouse, and the map is automatically saved upon exiting the program.

See Rover, following page

An X Window-based graphical interface has been developed to aid in using Rover.

Rover, cont. from page 14

Customizing

Users can bind arbitrary actions to mouse clicks or key clicks on nodes or links displayed on the map. These actions are defined in the user's .Xdefaults file. The example shown in Figure 4 is a user's .Xdefaults file which binds a "mouse button 1 down event" to opening an AIXterm window and pinging the node in question.

Figure 4

```
Nsfnet3*draw* XmPushButton.Translations: #augment \  
<Btn1Down>:      system( aixterm -e ping $1 ) \n\  
<Btn3Down>:      system( aixterm -e telnet $1 )
```

The "mouse button 3 down event" is likewise bound to opening an AIXterm window telneted into the node. The graphical display replaces all instances of \$1 with the IP Address of the node that was pressed.

The syntax for actions on links is similar. Merit is currently specifying the above tasks in Rover's .Xdefaults file by binding execution of a shell script rather than coding the actual command in the .Xdefaults. While originally done for readability, experience has shown that by executing a shell script, a layer of abstraction can easily be added. For instance, if there is a desire to get connected to a node, the shell script could try to use in-band access first, and if that failed, automatically use out-of-band access transparent to the operator.

One can imagine invoking a link testing shell script with a mouse or key click sequence that would perform link diagnosis. This shell script might query either end of the link for DSU and interface status, error counts, or other defined conditions.

Possible use with other tools

The Network Status File offers another advantage in that it can be used in other network management tools such as statistics collections. The status file provides a list of reachable nodes that is used during the collection of data from the network. As new nodes are discovered they are automatically added to the SNMP data collection list.

An additional feature is the automatic measuring of circuit quality by Merit's DSU circuit monitoring facility as new links are added.

Conceivably, this might permit automatic diagnosis of problems. Binding execution of a shell script to a node or link state transition would be one way of automating problem determination. This script could potentially update the operators' problem entry with the information that it discovered. Eventually, the problem might also be corrected automatically. In that case, the NOC would be informed of the occurrence and the corrective action taken.



Rover is continually being refined to stay abreast of the current needs of the Network Operations Center. For further information on this tool contact:

Bill Norton
Merit Network, Inc.
1071 Beal Ave.
Ann Arbor, MI 48109-2112
wbn@merit.edu

—Pat Smith, Merit/NSFNET

Users can bind arbitrary actions to mouse clicks or key clicks on nodes or links displayed on the map.

NSFNET Packet Traffic History



Merit/NSFNET
 2901 Hubbard, Pod G
 Ann Arbor, MI 48109-2016